

Inference and rare event simulation for stopped Markov processes via reverse-time sequential Monte Carlo

Jere Koskela

`koskela@math.tu-berlin.de`

Institut für Mathematik

Technische Universität Berlin

Straße des 17. Juni 136, 10623 Berlin, Germany

Dario Spanò

`d.spano@warwick.ac.uk`

Department of Statistics

University of Warwick

Coventry, CV4 7AL, UK

Paul A. Jenkins

`p.jenkins@warwick.ac.uk`

Departments of Statistics & Computer Science

University of Warwick

Coventry, CV4 7AL, UK

January 3, 2017

Abstract

We present a sequential Monte Carlo algorithm for Markov chain trajectories with proposals constructed in reverse time, which is advantageous when paths are conditioned to end in a rare set. The reverse time proposal distribution is constructed by approximating the ratio of Green’s functions in Nagasawa’s formula. Conditioning arguments can be used to interpret these ratios as low-dimensional conditional sampling distributions of some co-ordinates of the process given the others. Hence the difficulty in designing SMC proposals in high dimension is greatly reduced. We illustrate our method on estimating an overflow probability in a queueing model, the probability that a diffusion follows a narrowing corridor, and the initial location of an infection in an epidemic model on a network.

Keywords: intractable likelihood, rare event simulation, sequential Monte Carlo, stopped Markov process, time-reversal

1 Introduction

Sequential Monte Carlo (SMC) is a very general technique for sampling from a sequence of complicated distributions of increasing dimension and known pointwise up to a normalising constant. For a comprehensive introduction, see e.g. (Doucet et al., 2001; Liu, 2001; Doucet and Johansen, 2011). Briefly, a “cloud” of weighted particles is extended from one distribution to the next by a combination of sequential importance sampling and resampling. Each set of weighted particles then forms an empirical approximation of each subsequent distribution, provided adjacent distributions are sufficiently similar. This is typically the case when interest is in inference from a sequence of observations from e.g. a hidden Markov model, and thus SMC finds widespread use in the filtering literature (Doucet et al., 2001; Liu, 2001; Del Moral, 2004; Fearnhead, 2008; Doucet and Johansen, 2011).

In this paper we present a SMC algorithm for sampling trajectories of Markov chains started at the last exit time of an initial set I , and killed at the first hitting time of a terminal set T . By last exit time, we mean that trajectories of interest are not allowed to re-enter I before hitting T . Such Markov trajectories have a wide range of applications, such as population genetics (Griffiths

and Tavaré, 1994; Stephens and Donnelly, 2000; De Iorio and Griffiths, 2004a,b; Birkner and Blath, 2008; Hobolth et al., 2008; Griffiths et al., 2008; Birkner et al., 2011; Koskela et al., 2015), mathematical finance (Casella and Roberts, 2008), neuroscience (Bibbona and Ditlevsen, 2013), physics (Del Moral and Garnier, 2005; Johansen et al., 2006), and engineering (Blom et al., 2007; Lezaud et al., 2010). This problem is non-standard in SMC (Fearnhead, 2008) because the dimension of a particle can be random, and the intermediate target distribution of interest, namely that of a partially reconstructed chain conditioned on its eventual hitting of T before returning to I , is usually unavailable. In practice the intermediate target may be replaced with its unconditioned counterpart, in which case the current importance weight of a particle is in very poor correlation with its final weight. This renders the resampling step of a SMC algorithm counterproductive.

To circumvent this problem, our algorithm consists of proposing trajectories in *reverse* time, started from T and propagated until the first hitting time of I . The distribution of the reverse-time process is typically also intractable, but SMC techniques are used to mitigate the aforementioned problems, and produce properly weighted particle ensembles for unbiased estimators.

In practice, the utility of our method is in approximating quantities of the form $\mathbb{E}_\mu[f(X)]$, where f is an integrable function, X a Markov trajectory connecting I to T without returning to I , and μ is an initial distribution placing full mass on I . Dependency of f on the hitting time of T can also be incorporated without difficulty, as will be seen in later sections. The method is fairly general, but we can expect it to be particularly efficient in contexts where

- (i) the function f depends only on the terminal point in T , or a small length of trajectory preceding it,
- (ii) the support of $f|_T$ is small in terms of dimension and/or volume, while I is high dimensional and/or has large volume,
- (iii) the majority of the contribution to $\mathbb{E}_\mu[f(X)]$ arises from a region of low conditional probability given that the chain has hit T , and
- (iv) the process of interest is high-dimensional and transitions only alter a small number of components at a time.

Properties (i) and (ii) ensure that time-reversal is an effective strategy. In the extreme case where f is the indicator function of a singleton in T (corresponding to estimation of a conditional hitting density), I is a set which is hit by the reverse-time dynamics in finite time with probability 1, and T is a reverse-time entrance boundary, the optimal proposal distribution leading to zero variance estimators is the unconditional time-reversal. These conditions are very restrictive, but the fact that all of the examples in Section 4 violate at least one of them demonstrates that reverse-time proposals can still yield efficient algorithms under the milder conditions (i)-(iv).

Property (iii) is helpful in ensuring that T acts like a reverse-time entrance boundary with high probability, as proposal trajectories will naturally be pushed away from the rare hitting point of T and back towards I . Property (iv) means that it is only necessary to come up with a proposal distribution for the coordinates which differ between transitions, given the value of all other coordinates. This dimensionality reduction can greatly reduce the difficulty of designing proposal distributions in high dimension. Proposition 1 in Section 3 provides a precise formulation, and Sections 4.1 and 4.3 contain concrete examples.

Our method is reminiscent of existing forward-in-time SMC strategies known as multilevel SMC (Del Moral, 2004; Chen et al., 2005; Jasra et al., 2014), against which our method should be considered as a complementary rather than an alternative approach. Indeed, forward-in-time multilevel methods are well-suited for contexts involving the opposite to the point (ii) above, i.e. an initial set I of small volume and/or dimension, and a function f with support in T of

large volume and/or dimension.

While time-reversal has proved to be a successful tool for inference in population genetics (Griffiths and Tavaré, 1994; Stephens and Donnelly, 2000; De Iorio and Griffiths, 2004a,b; Birkner and Blath, 2008; Griffiths et al., 2008; Hobolth et al., 2008; Birkner et al., 2011; Koskela et al., 2015), and has also been used in rare event simulation (Frater et al., 1989; Anantharam et al., 1990; Frater et al., 1990; Shwartz and Weiss, 1993) and physics (Jarzynski, 2006), its use in combination with SMC has been limited to a few examples in population genetics (Chen et al., 2005; Jenkins, 2012; Koskela et al., 2015). It is in our opinion somewhat surprising that such an approach has not been combined with SMC for more general inference. Lin et al. (2010) use an auxiliary distribution operating in reverse time within SMC, but only to improve resampling, not to construct trajectories. Their time-reversal is also not constructed as an approximation of the reverse-time dynamics of their model. Jasra et al. (2014) use a coalescent-based example from population genetics to motivate their work, and highlight it as an example of more general SMC in reverse time, but still formulate their results forwards in time.

The examples presented in this paper are diverse: we consider an overflow probability in a queueing system, the probability that a diffusion trajectory remains contained in a narrow strip, and the initial location of an infection in an epidemic model on a network. These examples demonstrate the utility of reverse time proposals beyond the coalescent setting. The third example in particular would be challenging to solve with a forwards-in-time algorithm, as it is difficult to design a function whose level sets connect the initial condition to the desired rare event in a natural way (see (4) in the next section). Such a function is crucial for the implementation of forwards-in-time multilevel SMC algorithms, but not needed in our reverse time method.

The rest of the paper is laid out as follows. In Section 2 we formulate the rare event problem of interest and review forwards-in-time SMC. In Section 3 we present our reverse-time algorithm, and show how it naturally yields a dimension reduction in the SMC design task. Section 4 contains our example simulations, and Section 5 concludes with a discussion.

2 Problem formulation and SMC

Consider the canonical probability space

$$\left(\Omega := \prod_{n=0}^{\infty} E_n, \mathcal{F} := \bigotimes_{n=0}^{\infty} \mathcal{F}_n, \{X_n\}_{n=0}^{\infty}, \mathbb{P}_{\mu} \right), \quad (1)$$

where \mathbb{P}_{μ} is defined via its finite dimensional distributions as

$$\mathbb{P}_{\mu}(X_{0:n} \in dx_{0:n}) = \mu(dx_0) \prod_{i=0}^{n-1} P(x_i, dx_{i+1}). \quad (2)$$

Here $x_{0:n} := (x_0, \dots, x_n)$, and P is a given transition kernel. We assume both P and μ can be evaluated pointwise, but do not assume (1) is stationary or even has a stationary distribution. Following the convention of Del Moral (2004) (see page 51 in particular), we will also refer to (1) as the canonical Markov chain.

As in the previous section, let $I \subset \mathbb{N} \times \Omega$ be an initial set, and $T \subset \mathbb{N} \times \Omega$ be a space-time target set, which we assume has finite expected hitting time under the dynamics of $\{n, X_n\}_{n=0}^{\infty}$. We assume that $\mu(I) = 1$, and are interested in approximating functionals of trajectories

$$\mathbb{E}_{\mu}[f(\tau_T, X_{0:\tau_T}) | \tau_T < \tau_I] \quad (3)$$

for integrable functions f , where for a set $A \in \mathcal{B}(\mathbb{N}) \times \mathcal{B}(\Omega)$,

$$\tau_A := \inf\{n > 0 : (n, X_n) \in A\}$$

denotes the hitting time of A , and \mathbb{E}_μ denotes expectation with respect to \mathbb{P}_μ . We emphasize that these trajectories are defined between the *last exit time* of I and the *hitting time* of T . In particular, in our problem formulation trajectories of interest cannot re-enter I before hitting T . As an example, take T to depend only on space, and consider the hitting probability (resp. density) of a point $x \in T$ whenever Ω is discrete (resp. continuous), before any other point of T . The corresponding functional is

$$\mathbb{E}_\mu[f(\tau_T, X_{0:\tau_T}) | \tau_T < \tau_I] = \mathbb{E}_\mu[\mathbb{1}_{\{x\}}(X_{\tau_T}) | \tau_T < \tau_I].$$

As per point (iii) in Section 1, we assume that f places most of its mass on trajectories with terminal states that are unlikely under the dynamics of (1). Thus, approximating (3) can be seen as a rare event simulation problem of sampling trajectories with unlikely terminal states.

Sequential Monte Carlo is a standard method in rare event simulation (Rubino and Tuffin, 2009), and consists of constructing multiple trajectories in parallel by alternating between sequential importance sampling and resampling steps. Sequential importance sampling consists of sampling from a sequence of proposal distributions to build up a single, high-dimensional realisation. The proposals are typically not the conditional distributions of the model of interest, and so samples must be reweighted by the Radon-Nikodym derivative of the model and the proposal:

$$\begin{aligned} \mathbb{E}_\mu[f(X_{0:\tau_T})] &= \int f(x_{0:\tau_T}) \mathbb{P}_\mu(dx_{0:\tau_T}) = \int f(x_{0:\tau_T}) \frac{d\mathbb{P}_\mu}{d\mathbb{Q}_\eta}(x_{0:\tau_T}) \mathbb{Q}_\eta(dx_{0:\tau_T}) \\ &= \int f(x_{0:\tau_T}) \frac{\mu(x_0)}{\eta(x_0)} \prod_{n=0}^{\tau_T-1} \frac{P(x_n, x_{n+1})}{Q(x_n, x_{n+1})} Q(x_n, dx_{n+1}). \end{aligned}$$

Here \mathbb{Q}_η is a proposal distribution with initial law η and one-step transition kernel Q , and with $\mathbb{P}_\mu \ll \mathbb{Q}_\eta$. The resampling step is applied at intermediate times $t < \tau_T$, and consists of duplicating promising particles with high weight $\frac{\mu(x_0)}{\eta(x_0)} \prod_{n=0}^t \frac{P(x_n, x_{n+1})}{Q(x_n, x_{n+1})}$, while discarding particles with low weight.

Good choices of \mathbb{Q}_η and resampling schedule can dramatically reduce the variance of estimators, achieving the same asymptotic efficiency as the popular multilevel splitting method of rare event simulation under mild conditions Cérou et al. (2011). On the other hand, poor choices of \mathbb{Q}_η can yield estimators with higher variance than naive Monte Carlo (Glasserman and Wang, 1997). The optimal proposal is the law of the process conditioned on the event of interest, but computing this law involves the quantity of interest so the optimal algorithm is unimplementable in practice. The typical approach is to approximate the optimal proposal distribution using large deviations (Sadowsky and Bucklew, 1990).

Without resampling, the variance of estimators typically increases exponentially in the number of sequential steps even with well designed proposal distributions (Doucet et al., 2001; Liu, 2001). In the context of stopped processes, resampling should take into account the weight of the particle *and* the progress it has made towards the target set. This is achieved by introducing a sequence of intermediate sets, propagating all samples until they hit the next set, and perform resampling based on current weights once all particles have been stopped. This has been alternatively termed multilevel SMC or stopping time resampling in (Del Moral, 2004, Section 12.2) and (Chen et al., 2005), respectively. The good performance of our algorithm also relies on resampling, and hence should be regarded as reverse-time multilevel SMC.

Like the proposal distribution, the choice of resampling schedule also has a strong impact on the efficiency of the SMC algorithm. Few theoretical guidelines are available, though developments

have been made in determining good schedules adaptively both for multilevel splitting (C  rou et al., 2012) and SMC (Jasra et al., 2014). As with all multilevel splitting algorithms, the method of (C  rou et al., 2012) requires a reaction coordinate: a tractable function

$$\Psi : \Omega \mapsto \mathbb{R}, \quad (4)$$

such that both the rare event being estimated, as well as intermediate resampling stages, can be expressed as nested level sets of Ψ . The interested reader is directed to e.g. Section 2 of (C  rou et al., 2012) for details. Our method has the advantage of not requiring such a function. The examples in Section 4 will show that effective reaction coordinates are not always available, meaning that the advantage of not requiring one can be substantial. The particle MCMC method of Jasra et al. (2014) for designing efficient resampling schedules adaptively applies in our setting.

In summary, the main result of this paper is a general way to specify a SMC algorithm with a proposal distribution that constructs trajectories in reverse time. We will show how such proposals can be specified in very general settings, and that they yield good performance when combined with stopping time resampling. While performing stopping time resampling is crucial for the performance of our algorithm, we do not investigate optimisation of the resampling schedule. Our results apply to settings in which the conditions (i)-(iv) in Section 1 hold, and all three of the initial law μ , the forward transition density P and the reverse time proposal Q can be evaluated pointwise up to normalising constants.

3 Time-reversal as a SMC proposal distribution

In this section we review some relevant facts about time-reversal of Markov chains and introduce our generic SMC proposal distribution. Concrete examples can be found in Section 4.

We define the time-reversal of (1) by extending the chain to the negative time-axis, and letting

$$\left(\prod_{n=0}^{-\infty} E_n, \bigotimes_{n=0}^{-\infty} \mathcal{F}_n, \left\{ \tilde{X}_n \right\}_{n=0}^{-\infty}, \tilde{\mathbb{P}}_\nu \right)$$

denote the reverse-time chain. Note that the initial time-indices are set to 0 by convention, and are not necessarily intended to coincide with the starting time of (1). The law $\tilde{\mathbb{P}}_\nu$ is again defined via its finite dimensional distributions as

$$\tilde{\mathbb{P}}_\nu(dx_{0:-n}) = \nu(dx_0) \prod_{i=0}^{-n+1} \tilde{P}(x_i, dx_{i-1}).$$

Reverse time proposal distributions akin to (5) have been studied previously by Birkner et al. (2011) for certain population genetic models. The reverse transition kernel is related to its forward counterparts via Nagasawa’s formula (Rogers and Williams, 1994, Section VI.46):

$$\tilde{P}(x_i, x_{i-1}) = \frac{G(\mu, x_{i-1})}{G(\mu, x_i)} P(x_{i-1}, x_i), \quad (5)$$

where for a measurable set A ,

$$G(\mu, A) := \mathbb{E}_\mu \left[\sum_{i=0}^{\tau_T} \mathbb{1}_A(X_i) \right]$$

is the Green’s function of (1), and \mathbb{E}_μ denotes expectation with respect to \mathbb{P}_μ . When $A = \{z\}$ is a null set and the state space is continuous, we interpret the Green’s function $G(\mu, z)$ as

a density, which is assumed to exist. For simplicity we assume that all the transition kernels (resp. Green's functions) are absolutely continuous with respect to the same reference measure (e.g. Lebesgue or counting measure), and we will use the same notation for both transition kernels (resp. Green's functions) and their densities.

The Green's functions in (5) cannot be computed in most cases of interest, but their qualitative behaviour can often be described. We assume that such a description is available, and that it is possible to write down an approximating family of tractable functions with similar qualitative behaviour. It is not necessary for the match to be very precise, though better approximations yield more efficient SMC algorithms.

Our strategy for defining an SMC proposal is as follows:

1. Design an approximate Green's function $\hat{G}(\mu, x)$ to be substituted into (5) to yield an approximate reverse-time transition kernel \hat{P} and a proposal Markov chain

$$\left(\prod_{n=0}^{-\infty} E_n, \prod_{n=0}^{-\infty} \mathcal{F}_n, \left\{ \hat{X}_n \right\}_{n=0}^{-\infty}, \hat{\mathbb{P}}_\nu \right) \quad (6)$$

where $\hat{\mathbb{P}}_\nu$ is defined from its finite dimensional distributions via \hat{P} as before.

2. If necessary, modify $\hat{\mathbb{P}}_\nu$ locally to incorporate first hitting time constraints by preventing (6) from returning to T upon leaving it.
3. If necessary, further modify $\hat{\mathbb{P}}_\nu$ locally to ensure the expected hitting time of I by (6) is finite to ensure finite expected run time.

These steps can seem abstract because of their generality, but we will see in Section 4 that they can be carried out in many cases of interest.

Steps 2 and 3 could be incorporated automatically and more efficiently by considering the time-reversal of (1) after an appropriate Doob's h -transform (Rogers and Williams, 1994, Section VI.45), that is by considering forwards in time transition densities conditioned on hitting the rare event. However, the h -transform is typically intractable, whereas local modifications are widely implementable and still result in efficient algorithms when the dominant contribution to $\mathbb{E}_\mu[f(X_{0:\tau_T})]$ arises from a region of low \mathbb{P}_μ -probability and I lies in a region of high \mathbb{P}_μ -probability, as then the ratio of Green's functions will drive the process away from T and towards I without any conditioning (c.f. Property (iii) in Section 1).

Proposition 1 presents a practical way of designing approximate ratios of Green's functions for a wide class of models. For notational simplicity we assume a countable state space, but the same argument holds for continuous state spaces provided the Green's densities $g(x, z)$ exist.

Proposition 1. *Consider a transition of the Markov chain (1) from x_{n-1} to x_n , and suppose the state space can be partitioned so that $x_{n-1} = (z, y)$ and $x_n = (z, \bar{y})$. Assume that the conditional sampling distribution (CSD)*

$$\pi(y|z) := \mathbb{P}_\mu(Y_n = y | Z_n = z)$$

is independent of $n \in \mathbb{N}$ for \mathbb{P}_μ -almost every z . Then the ratio of Green's functions in (5) cancels to the ratio of CSDs:

$$\frac{G(\mu, (z, y))}{G(\mu, (z, \bar{y}))} = \frac{\pi(y|z)}{\pi(\bar{y}|z)}.$$

Proof. Let ∂ be a cemetery state, and define the process

$$X_t^\partial = \begin{cases} X_t & \text{if } t \leq \tau_T \\ \partial & \text{otherwise} \end{cases}.$$

Note that the laws of $\{X_n^\partial\}_{n=0}^{\tau_T}$ and $\{X_n\}_{n=0}^{\tau_T}$ coincide, and thus so do their Green's functions evaluated at states $(z, y) \in (T \cup \partial)^c$. Hence, for any such state, Fubini's theorem and conditioning on $Z_n = z$ yield

$$\begin{aligned} G(\mu, (z, y)) &= \mathbb{E}_\mu \left[\sum_{n=0}^{\infty} \mathbb{1}_{\{z, y\}}(Z_n^\partial, Y_n^\partial) \right] = \sum_{n=0}^{\infty} \mathbb{E}_\mu \left[\mathbb{1}_{\{z, y\}}(Z_n^\partial, Y_n^\partial) \right] \\ &= \sum_{n=0}^{\infty} \mathbb{E}_\mu \left[\mathbb{E}_\mu \left[\mathbb{1}_{\{y\}}(Y_n^\partial) | Z_n^\partial = z \right] \mathbb{1}_{\{z\}}(Z_n^\partial) \right], \end{aligned}$$

where $(Z_n^\partial, Y_n^\partial) := X_n^\partial$. Now

$$\pi(y|z) = \mathbb{E}_\mu \left[\mathbb{1}_{\{y\}}(Y_n^\partial) | Z_n^\partial = z \right]$$

is independent of n by assumption. Thus

$$G(\mu, (z, y)) = \pi(y|z) \sum_{n=0}^{\infty} \mathbb{E}_\mu \left[\mathbb{1}_{\{z\}}(Z_n) \mathbb{1}_{n:\infty}(\tau_T) \right] = \pi(y|z) \sum_{n=0}^{\infty} \sum_{t=0}^{\infty} \mathbb{P}_\mu(Z_n = z, \tau_T = t).$$

Now note that the final double sum cancels whenever the Green's functions are evaluated as ratios, which completes the proof. \square

Remark 1. The hypothesis of Proposition 1 is a weak conditional stationarity condition, and is relatively mild. However, since in practice we rely on ad hoc approximations to ratio of Green's functions, it is possible to extend the scope of the reverse-time framework by defining the proposal distribution based on a family of approximate CSDs $\hat{\pi}(y|z)$ even when Proposition 1 fails. In such a case, the ratio of approximate CSDs does not approximate a ratio of Green's functions, but can still yield an efficient algorithm if the qualitative behaviour of the approximate CSD resembles that of the true Green's function. For example, if $\pi(y|z)$ is not independent of time, but the time dependence in $\pi(y|z)$ is weak, then a CSD corresponding to a time homogeneous system may still yield a good approximation.

Because of their lower dimension, approximate CSDs are often much easier to design than either proposal kernels $\{Q(\cdot, \cdot)\}$ or approximate Green's functions $\{\hat{G}(\mu, \cdot)\}$. Indeed, we consider this dimensionality reduction in the design task one of the main advantages of the reverse-time framework.

Choosing a family of approximate CSDs $\{\hat{\pi}(\cdot|\cdot)\}$ and applying Proposition 1 to (5) yields proposal transition probabilities of the form

$$\hat{P}((z, \bar{y}), (z, y)) = \frac{\hat{\pi}(y|z)}{C((z, \bar{y}))\hat{\pi}(\bar{y}|z)} P((z, y), (z, \bar{y})),$$

where $C((z, \bar{y}))$ is a normalising constant, and the notation for the state variables is as in Proposition 1. The corresponding incremental importance weight at step n is

$$\frac{C((z, \bar{y}))\hat{\pi}(\bar{y}|z)}{\hat{\pi}(y|z)}.$$

Once a proposal chain has been constructed, functionals of interest can be unbiasedly estimated as

$$\begin{aligned} \mathbb{E}_\mu[f(X_{0:\tau_T})] &\approx \frac{1}{N} \sum_{j=1}^N f(x_{0:\tau_T}^{(j)}) \frac{d\mathbb{P}_\mu}{d\hat{\mathbb{P}}_\nu}(x_{0:\tau_T}^{(j)}) \\ &= \frac{1}{N} \sum_{j=1}^N f(x_{0:\tau_T}^{(j)}) \frac{\mu(x_0^{(j)})}{\nu(x_{\tau_T}^{(j)})} \prod_{n=1}^{\tau_T} \frac{\hat{\pi}(n(x_{n-1}^{(j)}, x_n^{(j)})|e(x_n^{(j)}, x_{n-1}^{(j)}))}{\hat{\pi}(n(x_n^{(j)}, x_{n-1}^{(j)})|e(x_n^{(j)}, x_{n-1}^{(j)}))} C(x_n^{(j)}), \end{aligned} \quad (7)$$

where $\{x_{0:\tau_T^{(j)}}^{(j)}\}_{j=1}^N$ is a sample from the SMC algorithm that uses (6) as its proposal mechanism, $e(x, y)$ is the vector consisting of those entries of x and y which are equal, and $n(x, y)$ is the vector consisting of the entries of x which are not equal to the corresponding entry of y . The purpose of this somewhat bulky notation is to encode the splitting of the state space into coordinates which remain constant between transitions, and others which change, that is introduced in Proposition 1. A formal algorithm specification for sampling properly weighted particle ensembles is given below in Algorithm 1.

Remark 2. Approximating (7) can be computationally daunting if $f|_T$ takes non-negligible values in a high-dimensional or large (in terms of Lebesgue-volume) subset of T , which is the rationale for property (ii) in Section 1. In such cases we do not expect the reverse-time approach to always be competitive with forwards-in-time algorithms, particularly if the initial set I is also small and hence difficult for the reverse-time chain to hit.

Remark 3. Normalising constants $C(x)$ in (7) would all be identically equal to one if an algorithm using the true ratio of Green’s functions could be implemented. Thus, the realised values of these constants for a given approximation could be used to design proposal distributions adaptively from trial runs, at least for discrete systems where the constants can be computed. We do not explore this possibility further in this paper.

We conclude this section with a discussion on the intuition of why a SMC proposal distribution of the form (6) should be efficient in the rare event simulation context outlined in Section 2. Loosely speaking, efficient rare event simulators for dynamical systems succeed by closely mimicking the conditional dynamics given the event of interest. In general, it is very difficult to identify good approximations to these dynamics, as they vary markedly from the unconditional dynamics of the system. In contrast, in reverse time the unconditional dynamics and the dynamics conditioned on a rare end point are very similar: both will rapidly leave the rare state, and return to regions of high probability. This behaviour is captured by the ratio of Green’s functions in (5), from which it is easy to see that steps that are proposed most often are those which, forwards in time, are both compatible with the dynamics of the model ($P(x_{i-1}, x_i) \gg 0$), and move the state towards the desired region of low probability ($G(\mu, x_{i-1}) > G(\mu, x_i)$). The reverse-time view permits initialising samples in the rare state of interest, and hence enables sampling trajectories that approximate the conditional dynamics given the rare event using only unconditional quantities — namely P and an approximation of $G(\mu, \cdot)$.

4 Numerical examples

In this section we present three simulation studies illustrating the reverse-time approach: the asynchronous transfer mode (ATM) network (Section 4.1), the hyperbolic diffusion (Section 4.2) and the susceptible-infected-susceptible (SIS) epidemic on a network (Section 4.3). All three examples fulfil conditions (i)-(iv) in Section 1, and consequently our reverse-time SMC outperforms a state-of-the-art adaptive multilevel splitting algorithm (C  rou et al., 2012) in the first two examples. Designing an efficient forwards-in-time SMC or splitting algorithm to tackle the third example would be a formidable task, and hence beyond the scope of this paper. In contrast, we show that designing an efficient algorithm in the reverse-time framework is straightforward.

4.1 The asynchronous transfer mode network

Our first example is the ATM network studied by Glasserman et al. (1999) in the context of rare events. The network consists of K sources, each of which is either on or off. Sources which are

Algorithm 1 Reverse-time multilevel SMC

Require: Particle number N , approximate CSD $\hat{\pi}$, stopping times $\{\tau_i\}_{i=n}^1$ such that $0 < \tau_n \leq$

```

 $\tau_{n-1} \leq \dots \leq \tau_1 = \tau_I.$ 
1: for  $j = 1$  to  $N$  do
2:    $X_0^j \sim \nu.$  ▷ Initial particle locations.
3:    $\bar{X}_0^j \leftarrow X_0^j.$  ▷ Duplicates used in resampling.
4:    $w_j \leftarrow \nu(X_0^j)^{-1}.$  ▷ Initial importance weights.
5:    $k_j \leftarrow 0, \bar{k}_j \leftarrow 0.$  ▷ Time indices.
6:    $A_j \leftarrow j.$  ▷ Ancestor indices.
7: end for
8: for  $i = n$  to  $1$  do ▷ Do for each level:
9:   if  $\text{ESS}(w_{1:N}) < N/2$  then ▷ Resampling check.
10:    for  $j = 1$  to  $N$  do
11:       $A_j \sim \sum_{k=1}^N w_k \delta_k.$ 
12:       $\bar{k}_j \leftarrow k_{A_j}.$ 
13:       $\bar{X}_{\bar{k}_j}^j \leftarrow X_{\bar{k}_j}^j.$ 
14:    end for
15:    Set  $\bar{w} \leftarrow N^{-1} \sum_{k=1}^N w_k.$ 
16:  else
17:    for  $j = 1$  to  $N$  do
18:       $A_j \leftarrow j.$ 
19:       $\bar{k}_j \leftarrow k_j.$ 
20:       $\bar{X}_{\bar{k}_j}^j \leftarrow X_{\bar{k}_j}^j.$ 
21:    end for
22:  end if
23:  for  $j = 1$  to  $N$  do
24:    if  $\bar{k}_j < \tau_i$  then ▷ First step in level.
25:       $k_j \leftarrow \bar{k}_j + 1.$ 
26:       $X_{k_j}^j \sim \frac{\hat{\pi}(n(\cdot, \bar{X}_{\bar{k}_j}^{A_j})|e(\cdot, \bar{X}_{\bar{k}_j}^{A_j}))P(\cdot, \bar{X}_{\bar{k}_j}^{A_j})\mathbb{1}_{T^c}(\cdot)}{\hat{\pi}(n(\bar{X}_{\bar{k}_j}^{A_j}, \cdot)|e(\cdot, \bar{X}_{\bar{k}_j}^{A_j}))C(\bar{X}_{\bar{k}_j}^{A_j})}.$ 
27:       $w_j \leftarrow \bar{w} \frac{\hat{\pi}(n(\bar{X}_{\bar{k}_j}^{A_j}, X_{k_j}^j)|e(X_{k_j}^j, \bar{X}_{\bar{k}_j}^{A_j}))C(\bar{X}_{\bar{k}_j}^{A_j})}{\hat{\pi}(n(X_{k_j}^j, \bar{X}_{\bar{k}_j}^{A_j})|e(X_{k_j}^j, \bar{X}_{\bar{k}_j}^{A_j}))}.$ 
28:       $A_j \leftarrow j.$ 
29:    end if
30:    while  $k_j < \tau_i$  do ▷ Propagate until next level.
31:       $X_{k_j}^j \sim \frac{\hat{\pi}(n(\cdot, X_{k_j}^j)|e(\cdot, X_{k_j}^j))P(\cdot, X_{k_j}^j)\mathbb{1}_{T^c}(\cdot)}{\hat{\pi}(n(X_{k_j}^j, \cdot)|e(\cdot, X_{k_j}^j))C(X_{k_j}^j)}.$ 
32:       $w_j \leftarrow w_j \frac{\hat{\pi}(n(X_{k_j}^j, X_{k_j+1}^j)|e(X_{k_j+1}^j, X_{k_j}^j))C(X_{k_j}^j)}{\hat{\pi}(n(X_{k_j+1}^j, X_{k_j}^j)|e(X_{k_j+1}^j, X_{k_j}^j))}.$ 
33:       $k_j \leftarrow k_j + 1.$ 
34:    end while
35:  end for
36: end for
37: for  $j = 1$  to  $N$  do
38:   Set  $w_j \leftarrow w_j \mu(X_{k_j}^j).$  ▷ Account for entrance law  $\mu$ .
39: end for

```

off do nothing, while sources which are on produce packets at rate λ . Packets are serviced by a common server with rate μ using the first-in-first-out policy. Off sources turn on at rate α_0 and

on sources turn off at rate α_1 . The state of the system is specified as $(i, j) \in \mathbb{N}_0 \times \{0, \dots, K\}$, where i denotes the number of packets in the queue and j the number of on sources.

Glasserman et al. (1999) estimated the probability of the queue length hitting a barrier $b \in \mathbb{N}$ before emptying, given an empty initial queue and $K\alpha_0/(\alpha_0 + \alpha_1)$ on sources. Reverse-time SMC could be used for this example by summing over all possible numbers of terminal on sources, but this results in a K -fold increase in computational burden and hence cannot be expected to be competitive with a forwards-in-time approach. We focus instead on the joint probability of an initially empty queue hitting a barrier b before emptying, with exactly k sources on at the hitting time, and assume the initial number of on sources is $\text{Bin}(K, \alpha_0/(\alpha_0 + \alpha_1))$ distributed. In this scenario a forwards-in-time algorithm would face the same difficulties as a reverse-time algorithm does in the scenario of Glasserman et al. (1999).

The initial set is $I = \cup_{r=0}^K \{(0, r)\}$, the target set is $T = \cup_{r=0}^K \{(0, r) \cup (b, r)\}$, the initial law is

$$\mu(\{(0, j)\}) = \binom{K}{j} \left(\frac{\alpha_0}{\alpha_0 + \alpha_1} \right)^j \left(\frac{\alpha_1}{\alpha_0 + \alpha_1} \right)^{K-j},$$

and the quantity of interest is the hitting probability

$$\mathbb{E}_\mu[f(\tau_T, X_{0:\tau_T})] = \mathbb{E}_\mu[\mathbb{1}_{\{(b,k)\}}(X_{\tau_T})].$$

The proposal is specified by defining approximate conditional distributions of i given j and j given i . These are denoted by $\hat{\pi}_i(i|j)$ and $\hat{\pi}_j(j|i)$ respectively, and we choose them to be

$$\begin{aligned} \hat{\pi}_i(i|j) &\propto \left(\frac{\lambda \max\{j, 1\}}{\mu} \right)^i \text{ for } i \in 0 : b \text{ and } j \in 0 : K, \\ \hat{\pi}_j(j|i) &\propto \binom{K}{j} \left(\frac{\alpha_0}{\alpha_0 + \alpha_1} \right)^j \left(\frac{\alpha_1}{\alpha_0 + \alpha_1} \right)^{K-j} \hat{\pi}_i(i|j) \text{ for } i \in 0 : b \text{ and } j \in 0 : K. \end{aligned}$$

The former is the true distribution of a queue with arrival rate λj and service rate μ whenever $\mu > \lambda j$, and well-defined otherwise as well because the range of possible values of i is finite. The minimum in $\hat{\pi}_i(i|j)$ is to allow for positive queue lengths even when no sources are active. We also employ stopping time resampling, with simulations stopped every time a new minimum queue length is reached in reverse time. Resampling takes place if the effective sample size is below half of the number of particles once all simulations have been stopped.

To present a comparison of our method, we also implemented the adaptive multilevel splitting algorithm of Cérou et al. (2012). This algorithm requires the user to specify a reaction coordinate Ψ (see (4) in Section 2), as well as a MCMC kernel targeting $\mathbb{P}_\mu(X_{0:\tau_T} | \Psi(X_{0:\tau_T}) \geq l)$ for any level l . We choose

$$\Psi(X_{0:\tau_T}) \equiv \Psi((i_t, j_t)_{t \in 0:\tau_T}) = \max_{t \in 0:\tau_T} \{i_t\}$$

as our reaction coordinate, and specify our MCMC kernel by

1. uniformly sampling a time point $t > 0$ along the trajectory,
2. uniformly sampling a new step from X_{t-1} to replace the current value of X_t ,
3. attaching the remaining trajectory $X_{(t+1):\tau_T}$ to the new value of X_t while keeping its step directions fixed.

Boundary handling was implemented by rejecting any proposals with number of on-sources outside the allowed interval $0 : N$, terminating a proposed trajectory upon its first hitting time of queue length 0 or b , and filling in incomplete trajectories from the unconditional ATM queue dynamics if a proposal yielded an incomplete trajectory that had not yet hit T . Proposals whose maximum queue lengths fell below the required level were also rejected. This adaptive algorithm

is known to lead to estimates with near-optimal variance among the class of multilevel splitting algorithms for a given Ψ (C  rou et al., 2012), though in practice the variance is also strongly impacted by the MCMC kernel.

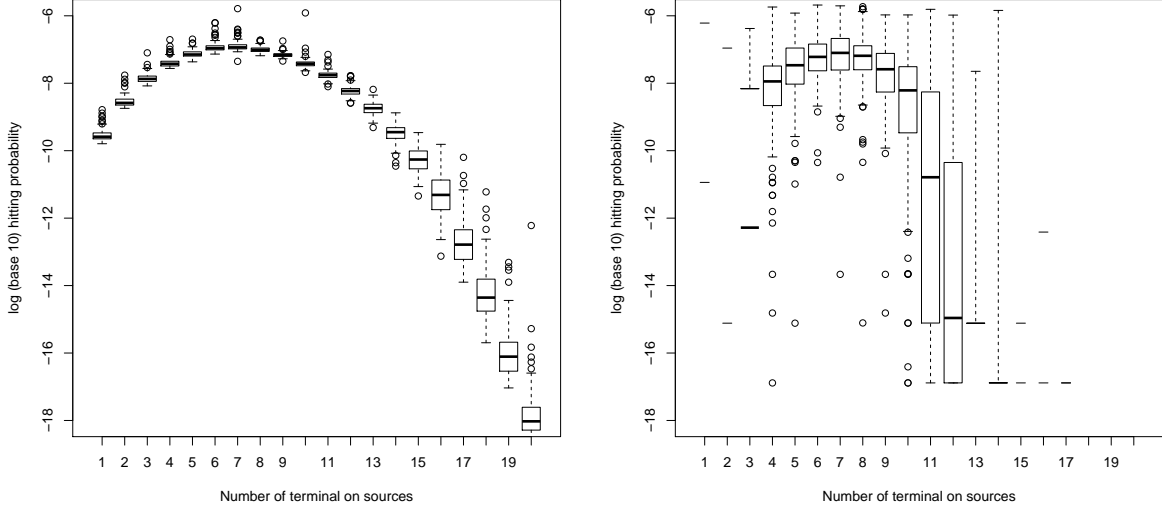


Figure 1: 100 replicates of simulated hitting probabilities using reverse time SMC (left) and adaptive multilevel splitting (right) on an ATM network with parameters $K = 20$, $b = 10$, $\lambda = 0.5$, $\mu = 10.0$, $\alpha_0 = 1.0$, $\alpha_1 = 3.0$. Particle numbers were 8 000 (SMC) and 10 000 (splitting) respectively, yielding comparable run times of 130 seconds per replicate each on an Intel i5-2520M 2.5 GHz processor. The missing and incomplete boxplots in the bottom figure correspond to terminal values of j that were hit by zero particles (resulting in a 0 estimate) in too many of the replicates for the default R boxplot method.

Figure 1 presents 100 replicates of simulated hitting probabilities of a queue length of 10 across all possible fixed numbers of terminal on sources. The SMC curve shows a high degree of smoothness, strongly suggesting the algorithm has converged, and the Monte Carlo variance is acceptably small for all but the rarest numbers of terminal on-sources. In contrast, the splitting algorithm has failed to produce estimates for the rarer terminal conditions at all, because too few of the particles being propagated forwards in time managed to hit these regions. The Monte Carlo variance of the estimates around the mode of the curve are also substantially larger than in the SMC case.

To conclude this section, Figure 2 shows 100 replicates of the reverse time SMC for a more challenging problem. In this case the adaptive multilevel splitting algorithm could no longer be run in feasible time due to both the length of trajectories and low acceptance probabilities resulting from the rarity of transitions towards the rare event. Both contributed to slow mixing of the MCMC step, resulting in infeasible run times. A more efficient MCMC kernel making use of block updates and a more finely tuned proposal mechanism could doubtless improve the performance of the multilevel splitting algorithm, but designing such a proposal is challenging without a detailed understanding of the dynamics of trajectories conditioned on the rare event. In contrast, reverse time SMC is readily implementable, and remains computationally feasible.

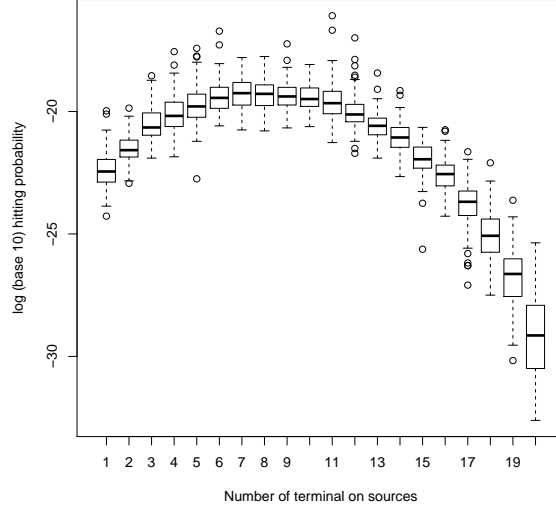


Figure 2: 100 replicates of simulated hitting probabilities using reverse time SMC on an ATM network with parameters $K = 20$, $b = 30$, $\lambda = 0.5$, $\mu = 10.0$, $\alpha_0 = 1.0$, $\alpha_1 = 3.0$, and 10 000 particles. The total run time was 22 hours, or around 40 seconds per independent simulation (20 terminal conditions \times 100 replicates) on an Intel i5-2520M 2.5 GHz processor.

4.2 The hyperbolic diffusion

The scalar hyperbolic diffusion is the solution of the SDE

$$dX_t = \frac{-X_t}{\sqrt{1+X_t^2}} dt + dW_t, \quad (8)$$

where $(W_t)_{t \geq 0}$ is a Brownian motion. It was introduced by Barndorff-Nielsen (1978) in connection to hyperbolic distributions in geostatistical modelling (Barndorff-Nielsen, 1977), and its heavier-than-normal tails have also made it a popular model in mathematical finance (Bibby and Sørensen, 2003).

The transition probabilities of the diffusion are intractable, but the stationary distribution is known to be the hyperbolic distribution

$$\pi(x) = \frac{1}{2K_1(1)} e^{-\sqrt{1+x^2}}, \quad (9)$$

where K_1 is the modified Bessel function of the second kind. We assume that the diffusion is started at stationarity, and focus on the probability that a trajectory lies in an interval (l_0, u_0) at time 0, and hits interval (l_t, u_t) at time $t \in \mathbb{N}$, without leaving the strip obtained by connecting l_0 to l_t and u_0 to u_t with straight lines at intermediate times. Similar containment probabilities have been studied e.g. by Casella and Roberts (2008) in the context of double barrier option pricing.

Formally, our inference problem is specified by the initial set $I = \{0\} \times (l_0, u_0)$, the target set

$$T = \bigcup_{s \in (0, t]} \left(\{s\} \times \left\{ \frac{l_t - l_0}{t} s + l_0, \frac{u_t - u_0}{t} s + u_0 \right\} \right)$$

the initial distribution $\mu = \pi$, and the quantity of interest

$$\mathbb{E}_\mu[f(X_{0:\tau_T})] = \mathbb{E}_\mu[\mathbf{1}_{\{t\}}(\tau_T)].$$

We consider a discretisation of (8), and use the Euler scheme with grid spacing Δ to define a family of approximate transition densities forwards in time:

$$P_\Delta((m, x), (n, y)) = \frac{\mathbb{1}_{\{m+\Delta\}}(n)}{\sqrt{2\pi\Delta}} \exp\left(-\frac{1}{2\Delta} \left[y - x \left\{1 - \frac{\Delta}{\sqrt{1+x^2}}\right\}\right]^2\right). \quad (10)$$

In this case (10) is also a Milstein scheme because of the unit diffusion coefficient.

We can use the discretised transition density (10) and the unconditional stationary distribution (9) to define a discretised reverse-time proposal:

$$\hat{P}_\Delta((n, y), (m, x)) \propto \frac{\pi(x)}{\pi(y)} P_\Delta((m, x), (n, y)) \mathbb{1}_{\left(\frac{l_t - l_0}{t} m + l_0, \frac{u_t - u_0}{t} m + u_0\right)}(x).$$

We also assume that Δ divides t exactly, and consider the analogous discretisation of the target set T . We neglect the issue of bias due to unobserved boundary crossings between time discretisation points, though more sophisticated interpolation schemes Gobet (2000) could also be implemented.

Note that (9) is not the stationary density of the discrete Euler approximation, which has different stationarity and reversibility properties than the continuous SDE (8). However, it can be expected to yield an efficient algorithm, because it is likely that (9) is a good approximation to the occupation measure of the Euler approximation, at least for a sufficiently small time step. We neglect the bias caused by using (9) as the initial condition of our discretised process.

The proposal distribution can be normalised numerically, and sampled by proposing $x(1 - \Delta(1 + x^2)^{-1/2})$ from a $\mathcal{N}(y, \Delta)$ proposal distribution, solving for x and accepting the proposal with probability $e^{1 - \sqrt{1+x^2}}$. We have incorporated step 2. of our strategy (see page 6) automatically in the above definition by rejecting proposed values outside the permitted strip. We also employ dynamic resampling, in which particles are resampled whenever the effective sample size falls below half the number of trajectories.

As a simulated example, we set $t = 2$, $(l_0, u_0) = (-1, 1)$, $\Delta = 0.01$ and $(l_t, u_t) = (5, 5.1)$. Note the wide initial condition, the narrow terminal condition and the upward slope of the strip, resulting in rare terminal conditions. All of these features make time reversal an attractive approach. Again, we compare our method to the adaptive multilevel splitting algorithm of Cérou et al. (2012), with reaction coordinate

$$\Psi(X_{0:\tau_T}) := \tau_T,$$

and MCMC kernel which, for level $\Psi(X_{0:\tau_T}) > l$

1. samples a non-initial time point s uniformly at random,
2. if $s \leq l$, then the proposal is $X'_s \sim U(l_s, u_s)$,
3. else $X'_s \sim P(X_{s-\Delta}, \cdot)$.

Again, proposed trajectories are terminated as soon as they hit the boundary of the target region, and incomplete proposals are propagated using P , the unconditional dynamics of the model, until the hitting time τ_T . Averaging over the initial distribution was done with naive Monte Carlo, with the overall containment probability taken as an average of containment probabilities from uniformly sampled initial conditions weighted by μ .

Figure 3 shows reverse-time SMC clearly outperforming adaptive multilevel splitting, mainly because the forwards-in-time splitting algorithm has to integrate over a much wider initial condition. Updating discretised diffusion trajectories using MCMC is also slow because of their length (200 steps in Figure 3) combined with the low acceptance probabilities associated with

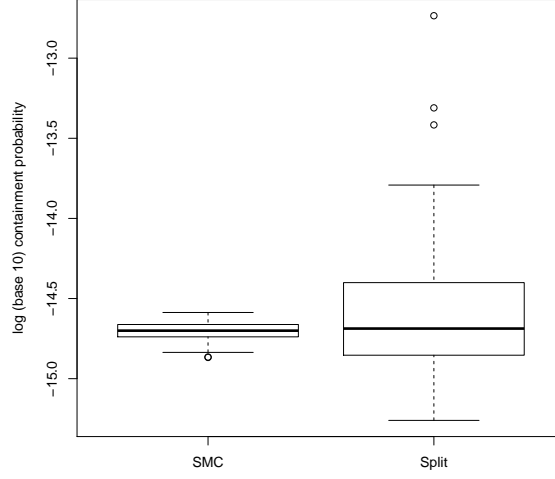


Figure 3: 100 replicates of simulated containment probabilities of the hyperbolic diffusion with initial interval $(l_0, u_0) = (-1, 1)$, terminal interval $(l_t, u_t) = (5, 5.1)$, trajectory length $t = 2$ and time step $\Delta = 0.01$. The particle numbers are 1 000 for SMC and 100 particles for each of 1 000 initial conditons for splitting, giving total run times of 7 seconds (SMC) and 240 seconds (splitting) per replicate on an Intel i5-2520M 2.5 GHz processor.

remaining contained in a narrowing interval. As in the previous section, Figure 4 demonstrates the performance of our SMC algorithm in a regime in which our multilevel splitting comparison was no longer feasible due to the increased trajectory length and resulting slower mixing.

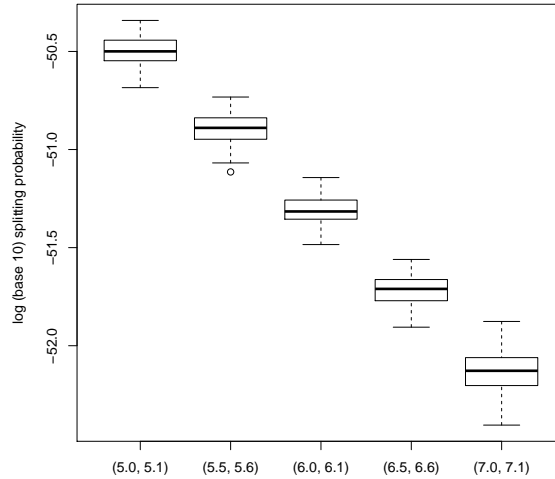


Figure 4: 100 replicates of simulated containment probabilities of the hyperbolic diffusion with initial interval $(l_0, u_0) = (-1, 1)$, terminal interval, trajectory length $t = 10$, time step $\Delta = 0.01$ and terminal interval as show. All estimates are using 1 000 particles with per replicate run times of 36, 39, 50, 55, and 66 seconds on an Intel i5-2520M 2.5 GHz processor, respectively.

4.3 The susceptible-infected-susceptible network

Consider a finite network with vertices V and undirected edges E , and with vertices labelled as either susceptible (S) or infected (J). For a vertex $v \in V$, let $l(v) \in \{J, S\}$ denote its label, $N_v := \{v' \in V : (v, v') \in E\}$ denote its neighbourhood and, for $a \in \{J, S\}$, let

$$N_v^a := \{v' \in V : (v, v') \in E \text{ and } l(v') = a\}$$

denote the sub-neighbourhood with label a . Then the susceptible-infected-susceptible (SIS) epidemic evolves as follows.

Every infected node is cured with rate $\beta > 0$, at which point it immediately becomes susceptible again. A susceptible node is infected by each infected neighbour at rate $\alpha > 0$, so that a vertex $v \in V$ becomes infected at total rate $\alpha|N_v^J|$. These types of dynamics on networks are popular models e.g. for the spread of biological epidemics in structured populations (Moore and Newman, 2000; Pastor-Satorras and Vespignani, 2001; Ganesh et al., 2015), malware in computer networks (Shah and Zaman, 2010), and rumours in social networks (Fuchs and Yu, 2015; Shah and Zaman, 2016), and are also sometimes referred to as contact processes. In addition, we impose a rate $\gamma > 0$ at which a new infection enters the network, infecting one uniformly sampled node. We assume further that new infections can only enter when all vertices are susceptible, i.e. only one infection can exist in the population at one time.

Suppose that there is no infection in the initial population, and that small infections go undetected. We define an infection as large once it infects a fixed number M of nodes, and assume that the labels of all nodes are immediately observed as soon as a large infection arises. Infection times are not observed, nor is any information about the history of the infection, such as whether a vertex that is now susceptible was previously infected. We are interested in inferring the initial location of the observed large infection, which may no longer be infected itself. Point estimators for similar inference problems have been studied in (Shah and Zaman, 2010; Fuchs and Yu, 2015; Shah and Zaman, 2016).

More formally we consider the jump skeleton of the above continuous time Markov process, let $l_t(v)$ denote the label of vertex $v \in V$ at time $t \in \mathbb{N}_0$, and let the Markov chain $\{X_t\}_{t=0}^{\tau_T}$ be given as

$$X_t = \{v \in V : l_t(v) = J\},$$

i.e. the set of infected vertices at time t . Then the initial condition I is the empty set, the target set T is the observed configuration of infected vertices, and the quantity of interest is

$$\mathbb{E}_\mu[f(X_{0:\tau_T}) | \tau_T < \tau_I] = \mathbb{E}_\mu[(\mathbb{1}_{\{v\}}(X_1))_{v \in V} | \tau_T < \tau_I],$$

where $(\mathbb{1}_{\{v\}}(\cdot))_{v \in V}$ denotes a vector of indicator functions indexed by vertices of the network. Note that approximating $\mathbb{E}_\mu[f(X_{0:\tau_T}) | \tau_T < \tau_I]$ using a forwards-in-time algorithm is challenging because it can be difficult to know a priori which nodes are likely initial infecteds, and hence the algorithm may spend much effort sampling trajectories of low probability. The problem also lacks a natural reaction coordinate (4); the candidate

$$\Psi(X_{0:\tau_T}) := \max_{t \in 0:\tau_T} \{|X_t|\}$$

succeeds in creating large epidemics, but does not favour infecting the observed nodes over any others. Hence, the probability of observing the correct infected configuration is still infeasibly low, even given the correct infection size. Because nodes can get uninfected and reinfected, the observed infection configuration cannot easily be expressed as the first hitting time of a region of the state space, making this problem very difficult to address. Neither of these problems causes any difficulty in reverse time: a reaction coordinate is not needed and the initial condition is integrated over automatically by the reverse time algorithm.

It remains to specify the proposal distribution, which we do by specifying a CSD of the label of one vertex given the labels of all the others. Conditioned on the labels of all other vertices, vertex $v \in V$ becomes infected at fixed rate $\alpha|N_v^J|$ and susceptible with rate β , so it is natural to base the CSD on the quantities

$$\left(\frac{\alpha|N_v^J|}{\alpha|N_v^J| + \beta}, \frac{\beta}{\alpha|N_v^J| + \beta} \right)$$

which are the true stationary probabilities of the label of $v \in V$ being J and S respectively, if the labels of all of its neighbours were fixed. However, in reverse time a vertex with no infected neighbours can also become infected, corresponding to a forwards-time infection spreading outwards and all connecting individuals becoming uninfected before an isolated leaf infection. Hence, we add a fixed $\varepsilon > 0$ to the infection rate of every vertex, yielding probabilities

$$\left(\frac{\alpha|N_v^J| + \varepsilon}{\alpha|N_v^J| + \beta + \varepsilon}, \frac{\beta}{\alpha|N_v^J| + \beta + \varepsilon} \right).$$

Finally, we also want infections to travel towards a common centre of mass. To modify our proposal distribution to achieve this, we require some notation. For a vertex v , let v_\uparrow denote the neighbour directly above, v_\downarrow the neighbour directly below, and analogously let v_\leftarrow and v_\rightarrow denote the two horizontal neighbours. Finally, let $d_\uparrow(v) = 1$ if the vertical coordinate of v is lower than the average vertical coordinate of infected vertices in the network, $d_\uparrow(v) = -1$ if the vertical coordinate of v is higher than the average among infected vertices, and $d_\uparrow(v) = 0$ if the two coincide. Define $d_{\leftrightarrow}(v)$ analogously in the horizontal direction. Now, let the CSD be given by

$$\hat{\pi}(l(v)|\{v', l(v')\}_{v' \neq v}) \propto \begin{cases} \frac{\alpha|N_v^J| + \varepsilon}{\alpha|N_v^J| + \beta + \varepsilon} w(v) & \text{if } l(v) = J \\ \frac{\beta}{\alpha|N_v^J| + \beta + \varepsilon} & \text{if } l(v) = S \end{cases},$$

where the weight $w(v)$ is defined as

$$w(v) = 2^{-[\mathbb{1}_J(l(v_\uparrow)) - \mathbb{1}_J(l(v_\downarrow))]d_\uparrow(v) - [\mathbb{1}_J(l(v_\rightarrow)) - \mathbb{1}_J(l(v_\leftarrow))]d_{\leftrightarrow}(v)}.$$

This weight is greater than one when a node that is proposed to be infected lies between an infected neighbour and the centre of mass of the infection, and less than one when a proposed infection is spreading from a neighbour away from the centre of mass. Hence, it instils a tendency for infections to grow towards a common centre of mass in reverse time. We note that an approximate CSD based on a larger neighbourhood size would yield a more accurate approximation at greater computational cost.

To study the Monte Carlo variance of our method, Figure 5 shows the maximum values of estimated likelihood surfaces across 100 iterations and various network sizes. The locations of the maximisers are summarised in Table 1. In all cases they concentrate on at most two neighbouring vertices with high probability, and invariably both are close to the truth. We note that there is no reason for the maximum likelihood estimator to coincide with the truth: the quantity of interest in Table 1 is the concentration of the maximum likelihood estimators across replicates. Likewise, there is no reason to expect any particular relationship between maximum likelihood values between network sizes. The purpose of Figure 5 is to demonstrate that reverse-time SMC naturally remains tightly concentrated on high probability regions of the path space even as the size of the state space grows.

To conclude this section, Figure 6 shows a representative likelihood surface for the initial location of the depicted observed epidemic. This relatively small network already illustrates that the distribution of the initial location is not homogeneous among observed infected locations, and that locations that are no longer infected cannot necessarily be ruled out.

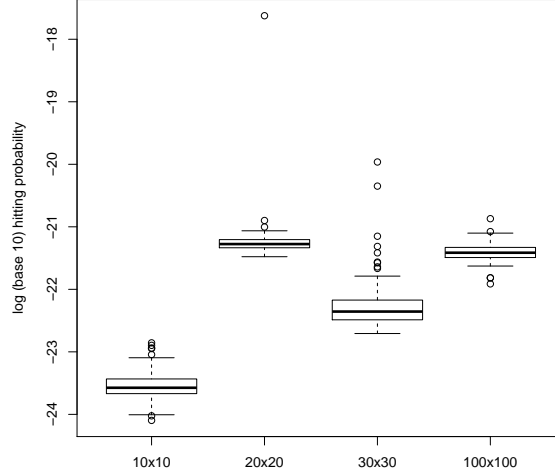


Figure 5: Distributions of the maximum likelihood value across 100 iterations as a function of grid size in a nearest neighbour network. The model parameters are $\alpha = 1$, $\beta = 12$, $\gamma = 1$, $\varepsilon = 10^{-4}$ for the 10×10 and 20×20 grids, $\varepsilon = 10^{-5}$ for the 30×30 grid and $\varepsilon = 10^{-6}$ for the 100×100 grid, so that $|V|\varepsilon \approx 10^{-2}$. Each simulation used 30 000 particles for per replicate run times of 40, 43, 44, and 89 seconds, respectively, on an Intel i5-2520M 2.5 GHz processor.

| 10×10 | 20×20 | 30×30 | 100×100 |
|----------------|----------------|----------------|------------------|
| (7, 3): 57 | (7, 8): 87 | (8, 7): 74 | (77, 50): 98 |
| (6, 3): 35 | (8, 7): 9 | (8, 8): 24 | (77, 49) : 1 |
| (5, 3): 6 | (7, 7): 3 | (8, 6): 1 | (77, 51) : 1 |
| (8, 3): 1 | (8, 8): 1 | (9, 7): 1 | True: (77, 49) |
| (7, 2): 1 | True: (7, 7) | True: (7, 6) | |
| True: (7, 2) | | | |

Table 1: Lists of locations and counts of maximum likelihood estimators of the initial infected position from the experiments described in Figure 5. Shown are the row and column labels with an $n \times n$ network represented as $(1 : n) \times (1 : n)$.

5 Discussion

We have presented a general framework for designing SMC proposal distributions which proceed backwards in time. Time-reversal makes it straightforward to ensure realisations of the process hit desired regions of the state space, essentially irrespective of the probability assigned to them by the law of the process of interest. Even the extreme case of conditioning paths on a terminal point of probability 0 can be dealt with easily. This makes time-reversal a natural and efficient choice when the end point of a path is known with high accuracy, but its initial distribution is flat. As most existing rare event and path simulation algorithms make the opposite assumptions about initial and terminal conditions, we expect time-reversal to be a useful tool in extending the scope of simulation-based inference and computation. Our simulation results support this expectation by demonstrating that simple reverse-time SMC schemes outperform a state-of-the-art, adaptive multilevel splitting scheme in settings similar to those described above, and remain feasible for problems for which multilevel splitting requires complex fine tuning.

Expressing the law of the reverse-time process via Nagasawa’s formula (5) often leads to a substantial reduction in the dimensionality of the design task of defining a proposal distribution

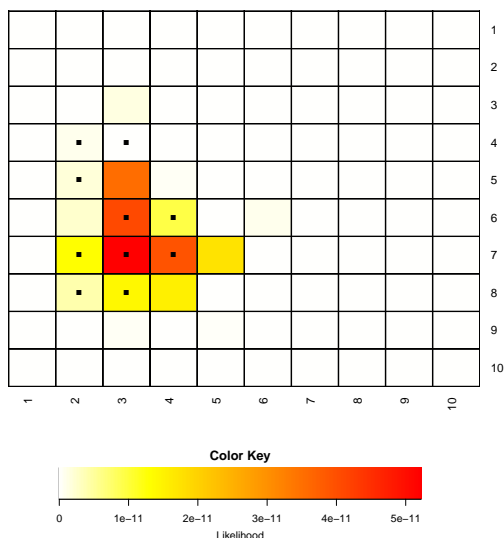


Figure 6: A representative simulated likelihood surface for the initial infection position in the 10×10 network. The model parameters are as described in Figure 5.

(c.f. Proposition 1). The difficulty of designing efficient proposal distributions in high dimension is a central barrier to practical SMC, so this cancellation of dimensions is an important advantage. Furthermore, we emphasize that it is not inherently linked to time-reversal: re-weighting jump probabilities by an approximate stationary distribution and cancelling out common coordinates would lead to a forwards-in-time proposal defined by low dimensional approximate CSDs. For rare terminal conditions a reverse-time approach is easier because the conditional and unconditional dynamics share the qualitative behaviour of rapidly leaving the rare state and moving towards a stationary mode. A forwards-in-time algorithm would have to use CSDs approximating the behaviour of an appropriate Doob's h -transform in order to drive the process away from modes and into the rare state. Nevertheless, we believe analogues of Proposition 1 to be a useful design tool beyond the scope of this paper.

All three example simulations considered in this paper have had the property that the proposal distribution could be normalised numerically, so that proposals could be sampled via standard methods. This property is computationally convenient, but is often not necessary since importance weights typically only need to be evaluated up to a normalising constant. Our reverse-time framework inherits this advantage from generic sequential Monte Carlo without added difficulty. Another advantages which reverse-time SMC inherits from standard SMC are that algorithms are straightforward to parallelise: use of GPUs for parallel Monte Carlo simulations has been found to speed up computations by up to 500 fold in comparison to the serial simulations shown in this paper (Lee et al., 2010). Further gains in efficiency can be made by reducing the required number of independent simulations through driving values (Griffiths and Tavaré, 1994) or bridge sampling (Meng and Wong, 1996). Such speed up, combined with the facts that

1. time reversal is advantageous in settings where forwards-in-time methods struggle, and
2. does not require a reaction coordinate (4) which can be difficult design in practice and is frequently necessary to implement forwards-in-time methods,

leads us to expect that a reverse time perspective can render many previously intractable problems amenable to practical computations.

Acknowledgements

The authors are grateful to Adam Johansen and Murray Pollock for fruitful conversations on rare event simulation, and to Ayalvadi Ganesh for pointing out the problem of inferring initial locations on networks. Jere Koskela was supported by EPSRC as part of the MASDOC DTC at the University of Warwick. Grant No. EP/HO23364/1. Paul Jenkins is supported in part by EPSRC grant EP/L018497/1.

References

- V. Anantharam, P. Heidelberg, and P. Tsoucas. Analysis of rare events in continuous time Markov chains via time reversal and fluid approximation. *IBM Research Report*, #RC 16280, 1990.
- O. Barndorff-Nielsen. Exponentially decreasing distributions for the logarithm of particle size. *Proc. R. Soc. London A*, 353:401–419, 1977.
- O. Barndorff-Nielsen. Hyperbolic distributions and distributions on hyperbolae. *Scand. J. Statist.*, 5:151–157, 1978.
- E. Bibbona and S. Ditlevsen. Estimation in discretely observed Markov processes killed at a threshold. *Scand. J. Stat.*, 40:274–293, 2013.
- B. M. Bibby and M. Sørensen. Hyperbolic processes in finance. In S. Rachev, editor, *Handbook of heavy tailed distributions in finance*, pages 211–248. Elsevier, Amsterdam, 2003.
- M. Birkner and J. Blath. Computing likelihoods for coalescents with multiple collisions in the infinitely many sites model. *J. Math. Biol.*, 57(3):435–463, 2008.
- M. Birkner, J. Blath, and M. Steinrück. Importance sampling for Lambda-coalescents in the infinitely many sites model. *Theor. Popln Biol.*, 79(4):155–173, 2011.
- H. A. P. Blom, G. J. Bakker, and J. Krystul. Probabilistic reachability analysis for large scale stochastic hybrid systems. In: *Proc. 46th IEEE Conf. Dec. Contr., New Orleans, USA*, 2007.
- B. Casella and G. O. Roberts. Exact Monte Carlo simulation of killed diffusions. *Adv. Appl. Probab.*, 40:273–291, 2008.
- F. Cérou, P. Del Moral, and A. Guyader. A non-asymptotic variance theorem for unnormalized Feynman-Kac particle models. *Ann. Inst. Henri Poincaré Probab. Stat.*, 47:629–649, 2011.
- F. Cérou, P. Del Moral, T. Furon, and A. Guyader. Sequential Monte Carlo for rare event estimation. *Stat. Comput.*, 22:795–808, 2012.
- Y. Chen, J. Xie, and J. S. Liu. Stopping-time resampling for sequential Monte Carlo methods. *J. R. Stat. Soc. B*, 67:199–217, 2005.
- M. De Iorio and R. C. Griffiths. Importance sampling on coalescent histories I. *Adv. in Appl. Probab.*, 36(2):417–433, 2004a.
- M. De Iorio and R. C. Griffiths. Importance sampling on coalescent histories II: Subdivided population models. *Adv. in Appl. Probab.*, 36(2):434–454, 2004b.
- P. Del Moral. *Feynman-Kac formulae: genealogical and interacting particle systems with applications*. Springer, New York, 2004.
- P. Del Moral and J. Garnier. Genealogical particle analysis of rare events. *Ann. Appl. Probab.*, 15:2496–2534, 2005.

- A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: fifteen years later. In D. Crisan and B. Rozovskii, editors, *The Oxford handbook of nonlinear filtering*. Oxford University Press, 2011.
- A. Doucet, J. F. G. De Freitas, and N. J. Gordon. *Sequential Monte Carlo methods in practice*. Springer, New York, 2001.
- P. Fearnhead. Computational methods for complex stochastic systems: a review of some alternatives to MCMC. *Stat. Comput.*, 18:151–171, 2008.
- M. R. Frater, R. A. Kennedy, and B. D. O. Anderson. Reverse-time modelling, optimal control and large deviations. *Syst. Control Lett.*, 12:351–356, 1989.
- M. R. Frater, R. R. Bitmean, R. A. Kennedy, and B. D. O. Anderson. Fast simulation of rare events using reverse-time models. *Comput. Networks ISDN*, 20:315–321, 1990.
- M. Fuchs and P.-D. Yu. Rumor source detection for rumor spreading on random increasing trees. *Electron. Commun. Probab.*, 20(Article 2):1–12, 2015.
- A. Ganesh, L. Massoulie, and D. Towsley. The effect of network topology on the spread of epidemics. In *Proc. 24th Annual Joint Conference of the IEEE Computer and Communication Societies (INFOCOM)*, volume 2, pages 1455–1466, 2015.
- P. Glasserman and Y. Wang. Counterexamples in importance sampling for large deviations probabilities. *Ann. Appl. Probab.*, 7(3):731–746, 1997.
- P. Glasserman, P. Heidelberger, P. Shahabuddin, and T. Zajic. Multi-level splitting for estimating rare event probabilities. *Oper. Res.*, 47:585–600, 1999.
- E. Gobet. Weak approximation of killed diffusions using Euler schemes. *Stoch. Proc. Appl.*, 87(2):167–197, 2000.
- R. C. Griffiths and S. Tavaré. Simulating probability distributions in the coalescent. *Theor. Popln Biol.*, 46:131–159, 1994.
- R. C. Griffiths, P. A. Jenkins, and Y. S. Song. Importance sampling and the two-locus model with subdivided population structure. *Adv. in Appl. Probab.*, 40:473–500, 2008.
- A. Hobolth, M. Uyenoyama, and C. Wiuf. Importance sampling for the infinite sites model. *Stat. Appl. Genet. Mol.*, 7:Article 32, 2008.
- C. Jarzynski. Rare events and the convergence of exponentially averaged work values. *Phys. Rev. E*, 73, 2006.
- A. Jasra, N. Kantas, and A. Persing. Bayesian parameter inference for partially observed stopped processes. *Stat Comput.*, 24:1–20, 2014.
- P. A. Jenkins. Stopping-time resampling and population genetic inference under coalescent models. *Stat. Appl. Genet. Mol. Biol.*, 11(1):Article 9, 2012.
- A. M. Johansen, P. Del Moral, and A. Doucet. Sequential Monte Carlo samplers for rare events. *Proceedings of the 6th international workshop on rare event simulation*, pages 256–267, 2006.
- J. Koskela, P. A. Jenkins, and D. Spanò. Computational inference beyond Kingman’s coalescent. *J. Appl. Probab.*, 52(2), 2015.
- A. Lee, C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes. On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *J. Comp. Graph. Stat.*, 19(4):769–789, 2010.

- P. Lezaud, J. Krystul, and F. Le Gland. Sampling per mode simulation for switching diffusions. *In: Proc. 8th Internk. Works. Rare-Event Simul., RESIM, Cambridge*, 2010.
- M. Lin, R. Chen, and P. Mykland. On generating Monte Carlo samples of continuous diffusion bridges. *JASA*, 105(490):820–838, 2010.
- J. S. Liu. *Monte Carlo strategies in scientific computing*. Springer, New York, 2001.
- X. L. Meng and W. H. Wong. Simulating ratios of normalizing constants via a simple identity: a theoretical exploration. *Statist. Sinica*, 6:831–860, 1996.
- C. Moore and M. E. J. Newman. Epidemics and percolation in small-world networks. *Phys. Rev. E*, 61:5678–5682, 2000.
- R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.*, 86:3200–3203, 2001.
- L. C. G Rogers and D. Williams. *Diffusions, Markov processes and martingales*, volume 1. Wiley, 2 edition, 1994.
- G. Rubino and B. Tuffin. *Rare event simulation using Monte Carlo methods*. Wiley, 2009.
- J. S. Sadowsky and J. A. Bucklew. On large deviation theory and asymptotically efficient Monte Carlo estimation. *IEEE Trans. Inf. Theory*, 36:579–588, 1990.
- D. Shah and T. Zaman. Detecting sources of computer viruses in networks: theory and experiment. *In Proc. ACM Sigmetrics*, volume 15, pages 5249–5262, 2010.
- D. Shah and T. Zaman. Finding rumor sources on random trees. *arXiv preprint*, 1110.6230, 2016.
- A. Shwartz and A. Weiss. Induced rare events: analysis via large deviations and time reversal. *Adv. Appl. Probab.*, 25(3), 1993.
- M. Stephens and P. Donnelly. Inference in molecular population genetics. *J. R. Statist. Soc. B*, 62(4):605–655, 2000.